LLNL-TR-407434

LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

WRF Test on IBM BG/L:Toward High Performance Application to Regional Climate Research

Hung-Neng S Chin

October 2, 2008

Disclaimer

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

WRF Test on IBM BG/L:
Toward High Performance Application to Regional Climate Research

Hung-Neng S. Chin

Atmospheric, Earth, and Energy Division
Lawrence Livermore National Laboratory
Livermore, CA 94551

1. INTRODUCTION

The effects of climate change will mostly be felt on local to regional scales (Solomon *et al.*, 2007). To develop better forecast skill in regional climate change, an integrated multi-scale modeling capability (i.e., a pair of global and regional climate models) becomes crucially important in understanding and preparing for the impacts of climate change on the temporal and spatial scales that are critical to California's and nation's future environmental quality and economical prosperity. Accurate knowledge of detailed local impact on the water management system from climate change requires a resolution of 1km or so. To this end, a high performance computing platform at the petascale appears to be an essential tool in providing such local scale information to formulate high quality adaptation strategies for local and regional climate change.

As a key component of this modeling system at LLNL, the Weather Research and Forecast (WRF) model is implemented and tested on the IBM BG/L machine. The objective of this study is to examine the scaling feature of WRF on BG/L for the optimal performance, and to assess the numerical accuracy of WRF solution on BG/L.

2. MODEL DESCRIPTION

The model used in this study is the latest version (Version 3.0.1) of the Advanced Research WRF (ARW) modeling system (http://www.mmm.ucar.edu/wrf/users/), a community model maintained by the National Center for Atmospheric Research (NCAR). The ARW is non-hydrostatic and fully compressible, and uses the sigma-pressure coordinate in the vertical axis to better simulate air flow over complex terrain. The model has a flux-form set of governing equations for better numerical conservation of mass and scalars. The ARW contains very complete model physics, and multiple options for each physical process, such as cumulus convection, microphysics of cloud and precipitation, long- and shortwave (LW and SW) radiation, turbulence and diffusion, planetary boundary layer (PBL), surface layer, and soil layer representations. The reader is referred to Skamarock et al. (2007) for further details.

All WRF simulations on BG/L shown in this study use the default mode (both CPUs per node used for computation and signal communication), which is run slightly faster than the virtual model (one CPU for computation and the other for signal communication). The third order Runge-Kutta scheme is used in the time-splitting integration with sound waves treated explicitly in the horizontally and implicitly in the vertical on shorter sub-steps, and 5th and 3rd order scheme for the horizontal and vertical advection, respectively.

We present two types of WRF simulations in this work; (1) idealized simulations and (2) real-data simulations (See Table I). Both types of simulations have the domain top residing at the altitude of 50 mb (~ 20 km). The former type uses a single domain to simulate a mid-latitude severe storm for testing the scaling feature of WRF on BG/L and demonstrating the resolution (1-km, 2-km, and 4-km) impact on the model solution. All idealized simulations contain 40 grid points in the vertical. The physics parameterizations used in idealized simulations include the Goddard microphysics scheme and the Monin-Obukhov surface scheme. The Goddard microphysics used in this study is based on its default setting for the 3ice scheme with graupel (Tao et al., 2003a).

Table I. Experiment outline.

Idealized			Real-data		
Simulation Type	Resolution	Dimension (X, Y)	Nested Domain	Resolution	Dimension (X, Y)
Scaling Test	2-km	(100, 100)	Outer	36-km	(90, 120)
Resolution Impact	1-km	(800, 800)			
	2-km	(400, 400)	Inner	12-km	(100, 130)
	4-km	(200, 200)			

The latter one is performed with boundary conditions from LLNL's 1.25° x 1° global climate model. The real-data simulations shown in this study are conducted with two levels of nested grids, which are in two-way coupling. The outer coarser-grid domain has a grid spacing of 36 km with a grid-size ratio of 3 to define the inner fine-grid domain. The inner grid (Δx = 12 km) is chosen to better resolve the local topography of the coastal ranges. A total of 100 (90) and 130 (120) grid points are selected to define the east-west and north-south axis of the finer (coarser) grid domain, respectively. Only the results from the 12-km resolution of the inner grid domain are shown in this study. The vertical axis contains 31 levels with 15 m resolution near the ground and gradually coarser resolution aloft. The 30-second resolution static field data (land-use, terrain, and soil-type) are used to initialize simulations. The real-data simulations use more physics than its idealized counterpart. These physical processes include Grell-Devenyi cumulus scheme (Grell and Devenyi, 2002), Goddard microphysics parameterization (Tao et al., 2003a), Rapid Radiative Transfer Model (RRTM) longwave radiation (Mlawer et al., 1997), Dudhia (1989) shortwave radiation, Yonsei University (YSU) boundary layer scheme (Hong et al., 2006), and Rapid Update Cycle (RUC) surface parameterization (Smirnova et al., 2000).

3. RESULTS

a. Scaling Test

The scaling test of the idealized simulation indicates that the computational efficiency of WRF simulation on BG/L degrades rapidly when the decomposed domain size per CPU is smaller than 25 x 25 grid points (i.e., 16 CPUs used) as a result of increasing overhead on MPI communication (Fig. 1). Therefore, this decomposed domain size is used as an optimal setting for WRF simulations on BG/L. The simulation is forced to terminate due to poor computational efficiency as the total CPU used reaches more than 128. Therefore, the minimum number of grid points per CPU that can run on BG/L is 81 (9 x 9 grid points per CPU). As seen in Fig. 1, CPU time decreases with the increasing number of CPUs used. However, total CPU time with a higher number of CPUs than the optimal size is much larger than its counterpart that contains more than 25 x 25 grid points per CPU as a result of increasing overhead on MPI communication.

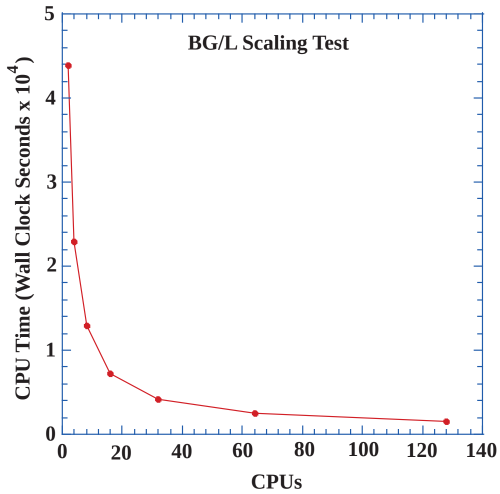

Fig. 1. The scaling test of WRF idealized simulations with a grid resolution of 2-km. The horizontal grid dimensions are 100 x 100 grid points

b. Resolution Impact of Idealized Simulations

To assess the impact of horizontal resolutions (1, 2, and 4-km) on the model performance, an idealized case of a mid-latitude squall-line storm (i.e., a standard WRF benchmark simulation) is selected for this test. The model domain is 800-km wide in the horizontal and 20-km deep (40 layers) in the vertical. During the transition stage of BG/L to the open computing environment, a maximum of 1024 CPUs is allowed for our test. A 3-D structure (in terms of the iso-surface of cloud concentration) of the simulated severe storm in a resolution of 1-km at 6-h simulation time is displayed in Fig. 2. The total CPU time for this high resolution 6-h run is 4.3 hours (wall clock). The anticipated surface gust front and upper-level (near 10 km) anvil cloud are well resolved in this simulation. Similar structure is also captured in coarser resolutions of 2 and 4-km (not shown). The impact of grid resolutions on the model solution is not clearly identified in the 3-D cloud envelope, but this impact is well captured within the simulated storm. The finer grid resolution resolves more detailed structure within the storm (Fig. 3).

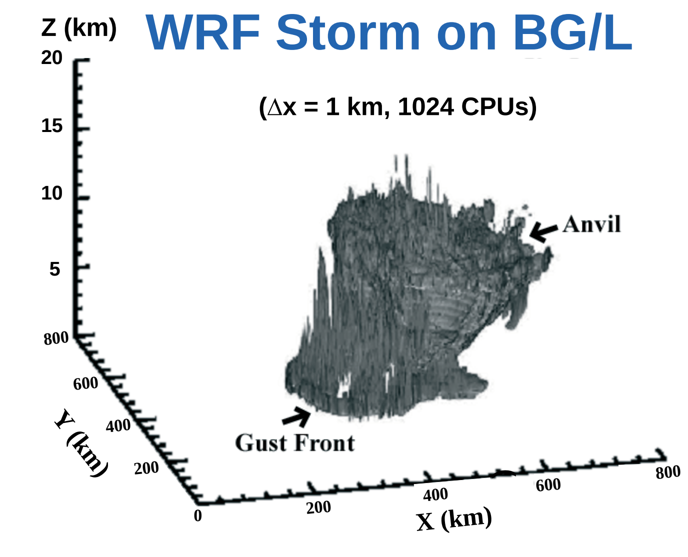

Fig. 2. The 3-D structure of simulated squall-line storm with the resolution of 1 km at 6-h of simulation time.

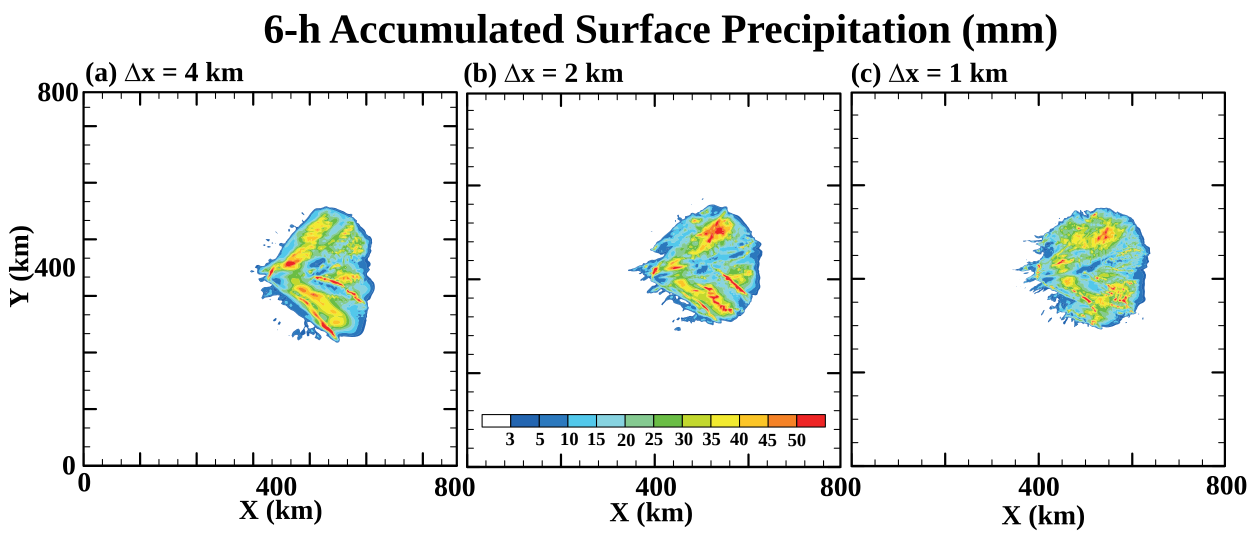

Fig. 3. 6-h accumulated surface precipitation. (a) Δx = 4 km. (b) Δx = 2 km. (c) Δx = 1 km.

c. Regional Climate Simulations

The WRF model has been extensively tested and validated on ZEUS in LLNL's regional climate research. Therefore, the WRF solution on ZEUS is used as a benchmark run to evaluate the WRF performance on BG/L. Monthly WRF simulations using the boundary conditions from LLNL's CCSM January data over 10 consecutive years are conducted for this purpose. Results from monthly averaged surface precipitation rate indicate that the discrepancy between BG/L and Zeus is very small (Fig. 4). The horizontal area average of surface precipitation differs only 0.4% between these two 10-Y averaged runs. Within these 10 January simulations, the maximum differences of surface precipitation between single January Zeus and BG/L runs are 1.78% and 0.98%, and the rest range from 0.03% to 0.51%. This test demonstrates that BG/L is able to produce reliable solution for WRF simulations.

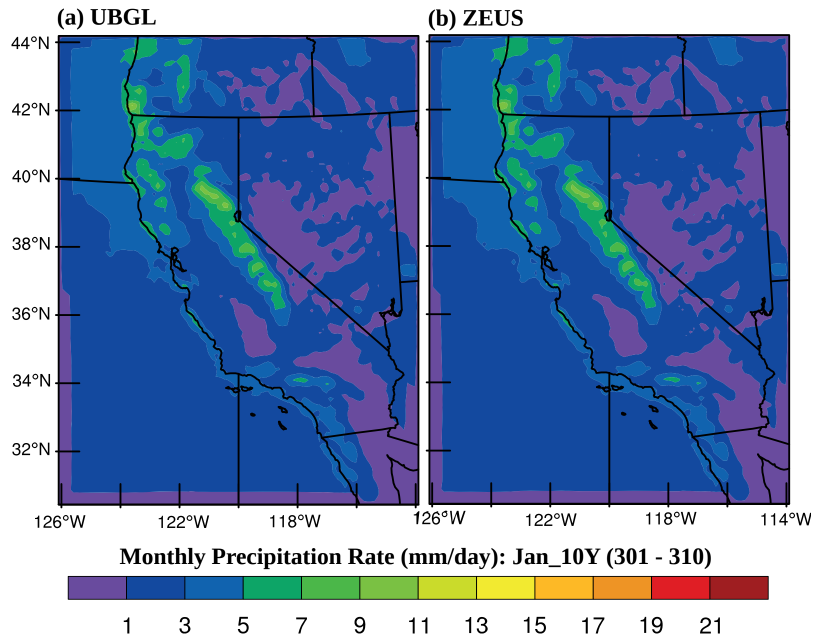

Fig. 4. Monthly surface precipitation rate (mm day^{-1}) of January simulations over 10 years of the CCSM present-climate case.

Caveat should be noted that the low memory feature of BG/L substantially slows down WRF real-data simulations since more memory is needed for multiple nests. For this benchmark test with the same number of CPUs, the CPU time used on BG/L is 9 times more than its ZEUS counterpart. Note that the CPU chip speed is 2.4 GH on Zeus and 0.7 GH on BG/L. This memory factor is expected to increase if more nested domains are required for finer resolution applications. Although BG/L can provide more CPUs to obtain memory, the extra overhead on MPI communication appears to offset this BG/L strength. Therefore, more tests are recommended to determine an optimal configuration of WRF's real-data simulations for the petascale applications.

4. SUMMARY AND DISCUSSION

The WRF model has been successfully implemented on BG/L, and evaluated against its counterpart on ZEUS. Preliminary tests indicate that BG/L is able to produce reliable solution for WRF simulations. Nonetheless, the low memory feature of BG/L raises an extra computational barrier for WRF's finer resolution applications in climate change research. In particular, 1-km resolution information is highly useful for an accurate water management system. Based on our scaling test, we estimate to have 8 times of BG/L total processors to simulate 5 years per (CPU time) day for a WRF 1-km simulation over California. This number does not include the memory slowdown factor. A much larger computer power would be expected for the need in the nation wide simulations. Therefore, more research is recommended to optimize the BG/L usage for petascale applications to climate change research.

Acknowledgments. The author wishes to thank Tom Spelce and Bor Chan at Livermore Computing for their help during the implementation of WRF on the IBM BG/L machine, and Julie Lundquist for valuable suggestion to improve the draft. The author also thanks Doug Rotman and Dave Bader for their support on this project. This work is performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

References

Dudhia, J., 1989: Numerical study of convection observed during the Winter Monsoon Experiment using a Mesoscale Two-Dimensional Model. *J. Atmos. Sci*., **46**, 3077-3107.

Grell, G.A. and D. Devenyi, 2002: A generalized approach to parameterizing convection combining ensemble and data assimilation techniques, *Geoph. Res. Let*., **29**, NO 14., 10.1029/2002GL015311, 2002.

Hong, S.-Y., Y. Noh, and J. Dudhia, 2006: A new vertical diffusion package with an explicit treatment of entrainment. *Mon. Wea. Rev*., **134**, 2318-2341.

Mlawer, E., S. Taubman, P. Brown, M. Iacono, and S. Clough, 1997: Radiative transfer for inhomogeneous atmospheres: RRTM, a validated correlated-k model for the longwave. *J. Geophys. Res.*, **102**, 16663-16682.

Skamarock W.C., and Coauthors, 2007: A Description of the Advanced Research WRF Version 2. NCAR TECHNICAL NOTE, NCAR/TN–468+STR.

Smirnova, T., J. Brown, S. Benjamin, and D. Kim, 2000: Parameterization of cold-season processes in the MAPS land-surface scheme, *J. Geophys Res*., **105**, 4077-4086.

Solomon, S., D. Qin, M. Manning, Z. Chen, M. Marquis, K. Averyt, M. Tignor, and H. Miller, eds., 2007: Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, 996 pp.

Tao, W.-K., J. Simpson, D. Baker, S. Braun, M.-D. Chou, B. Ferrier, D. Johnson, A. Khain, S. Lang, B. Lynn, C.-L. Shie, D. Starr, C.-H. Sui, Y. Wang and P. Wetzel, 2003a: Microphysics, radiation and surface processes in the Goddard Cumulus Ensemble (GCE) model, *A Special Issue on Non-hydrostatic Mesoscale Modeling, Meteorology and Atmospheric Physics*, **82**, 97-137.